FABIEN BETREMIEUX:	Good morning, good afternoon, and good evening, everyone. Welcome to our meeting of the Thick WHOIS IRT on Tuesday the 14th of June, 2016. My name is Fabien Betremieux with the Global Domains Division of ICANN.
	I see that we have from the IRT today with us Jody, Joyce, Roger, Steve, and Theo. I hope that other members will join us when we get started. Am I missing anyone, maybe, on the line?
DENNIS CHANG:	Hello, Fabien. This is Dennis Chang. May I make a quick announcement?
FABIEN BETRIEMIEUX:	Yes, Dennis. Please.
DENNIS CHANG:	Hi. This is Dennis Chang from ICANN staff. I would like to make a quick announcement on Fabien's other project. The project is called Alice. I'm happy to announce that Fabien has a new daughter [inaudible] birth yesterday named Alice. So congratulations are in order for his delivery ahead of schedule. Three weeks, right, Fabien?
FABIEN BETRIEMIEUX:	Yeah, absolutely. Thank you, Dennis. This was quite a surprise, but we are certainly blessed and we're happy. Thank you all.

UNIDENTIFIED MALE: Congratulations.

FABIEN BETREMIEUX: Thank you. Thank you very much. Let me go through a few reminders before we start with our agenda today. Make sure that your line is muted when you're not speaking. The meeting is recorded and will be transcribed. For the purpose of the transcript, please don't forget to say your name when you speak. Finally, if any time you'd like to get into the queue, please do so by raising your hand in the Adobe Connect room.

All right. For our agenda for today, what I'm suggesting we do is that we work off of our scorecard, as we started doing last week. We will get some of our priority items that are open and pretty important and on which we might have some updates. So we'll look at everything – registration track. We'll go as well to the new registration track and then look at these priority items. If we have more time, we'll look at the other various items that we have on the scorecard. Finally, we'll quickly look at our next meeting.

Any comments, questions, or suggestions on the agenda?

Hearing none. All right, let me load the scorecard. By the way, I've sent earlier today the updated scorecard, so you should have both a clean and a redline version in your e-mail. That was sent to the mailing list. If you'd like to open that up to follow on your [inaudible], please do so.

I marked the items that I think are priority items for today with a blue background in the number column, in the first column of the able. What I propose we do is we look at the validation rules for a start.

As you may remember, last week we initiated sharing the rules that have been discussed so far with the Registrar Stakeholder Group and the Registry Stakeholder Group. In our table, it's 3C.

Let me stop here and see if there is any update at this point that we can start discussing, or whether we need to wait a bit more. Let me stop here and see if you and potentially Mark may be able to provide some update on what they've heard from the stakeholder groups.

Theo, please, go ahead.

THEO GEURTS: Thank you, Fabien. I will take it off [this day]. 3B [inaudible] some input there, but after re-reading it and looking at it again, I think you can strike my name there. Thanks.

FABIEN BETREMIEUX: Okay. Theo, just to be clear, we have left open 3B, which is regarding the potential incentives or other measures that would provide an incentive for registrars to make sure they provide all the data that's available. It seems that there's no such suggestion at this point, so we would be closing that item.

Is anybody opposed to doing that?

I'm not seeing any objection marks or anybody raising their hand in the Adobe Connect, so we'll just close 3B. We will base our work off the principles that are agreed on in 3B, which is that the goal is that the RDDS output be the same before and after a transition, whether the validation was [inaudible].

EN

Okay. Let's then move on. Theo, on 3C, do you have any inputs from initial discussions in the Registrar Stakeholder group you could share with the IRT?

THEO GEURTS: Yes, Fabien. Just a quick update. It's just really what it was on the mailing list. We've had some discussion with the Registrar Stakeholder Group. There were no strong – well, I couldn't sense any rejections regarding the validations rules there. So we are good to go there, and that was the initial question there.

There was some input that I shared with the group here, and I noticed that you put out some modifications there in the other document, Fabien, just earlier on. I saw you incorporated them all already in the document. We can discuss them later.

FABIEN BETREMIEUX: I believe that's really the bulk transfer, right? Is that what you're talking about?

THEO GEURTS: Yeah, the bulk transfer, the alternative method there. It seems that some registrars are still looking at the RDE deposits as an option there to make things easier for them if they don't want to follow the EPP path. But we still need to flesh that out, if that is even a possibility. I'm not sure there. I just thought it mentioned it a couple of times before I pass it on to the group here.

FABIEN BETREMIEUX: Thank you, Theo. We'll get to that once we're done with discussing the validation rules.

Before we close the discussion on the validation rules, let me just turn to make Mark to see if there is any update on the registry side.

MARK MONITOR: Hi, Fabien. Thank you.

FABIEN BETREMIEUX: Hello, Mark.

MARK MONITOR: That was a good one. Sorry. Congratulations. Did I hear that your child was born?

FABIEN BETREMIEUX: Yeah. Thank you very much.

MARK MONITOR: Congratulations.

FABIEN BETREMIEUX: Thank you. Thank you so much. She was born three weeks in advance, so that was a bit of a surprise.

EN

MARK MONITOR:
Ah, all right. And you still made the call. I appreciate your dedication there.

In the validation rules, the Registry Stakeholder Group is meeting tomorrow. I was hoping to get on the agenda for that, but I wasn't able to. They have a full agenda for tomorrow. So what I'll do is I'll post it to their e-mail discussion and request feedback from that.

I did get your overview or I guess your intro to it, so thank you for that. I think that's very good and very helpful. So I'll just take that and use that in my e-mail to the rest of the Registry Stakeholder Group.

I also want to add that I don't anticipate that there'll be any issues or objections to that. I think it's pretty straightforward. But, again, I think it just makes sense to update the Registry Stakeholder Group on that and give them a chance to provide feedback if they have any. So I'll go ahead and do that, and I'll provide a further update at next week's meeting.

Thank you.

FABIEN BETREMIEUX:
Thank you, Mark. Theo, I see your back in the queue. Is that correct?

THEO GEURTS:
Yeah. That's correct. New hand. Just to point out – this is probably not on the list, but I'm going to point it out again – we as registrars are going to have a policy call next Monday, so there might be some information that we get from there. I don't think we get many or anything at all, but it will be discussed. So, we'll see. Thanks.

EN

FABIEN BETREMIEUX: Thank you, Theo. Before we move on, I just want to make my own announcement, and that's for Theo's election, I understand, as the Secretary of the Registrar Stakeholder Group. Congratulations to you, Theo.

THEO GEURTS: Yeah. Well, we still need to run the elections, so we're not there yet. We have a [inaudible]

FABIEN BETREMIEUX: Sorry. The nomination. Apologies.

THEO GEURTS: Yeah. We have the nomination. Yeah, we still have the elections going on. The nominations were already posted, and there were many.

FABIEN BETREMIEUX: Yeah. My apologies for the confusion, Theo. But congratulations for your nomination.

THEO GEURTS: Thanks.

FABIEN BETREMIEUX: One last thing on validation rules. I think, for 3D, we discussed it in our last meeting and there seems to have been an understanding that there

EN

would be no difference once the backfill of existing registrations are complete – that is, once all this existing registration data is transferred – there would be no difference made between existing and new registration in terms of what validation rules would apply. There was no objection to that proposal. I left it ongoing in the table. But it struck to me that this is something we may want to close.

So I'm talking about 3D here and whether there needs to be a difference made in terms of treatment between new and existing registration, in terms of validation rules, once the transfer of registrations are complete.

Let me see if there's anybody opposed to closing that one based on the conversation we had last week, which I put in the document.

Mark?

MARK MONITOR: Thank you, Fabien. Not an objection, just making sure we're all on the same page here. I think, basically, the proposal is that the registry would operate the minimum validation rules to start with, right? This would be to aid in the backfill or existing registration contact data. But those minimum validation rules would apply to new and existing registrations.

The end date, or when we would cut over from the minimum validation rules to what I'll call standard validation rules, would occur at the end date of the backfill period; so whatever timeline we define as the end of the backfill period, which may or may not overlap with the new registration transition period. But whatever we define as the end date

for the backfill period is when we would cut over from the minimum validation rules to the standard validation rules.

I think that's what we're all agreeing to. I think that makes sense from my perspective, but I just wanted to restate it to make sure we're all on the same page. Thank you.

FABIEN BETREMIEUX: Thank you, Mark. Would anybody like to share any comments or objections to closing 3D?

Okay. I'm not hearing anyone. I'm not seeing any hands raised, nor any [inaudible] in the Adobe Connect, so we'll consider that one closed for now. If at any point in time somebody would like to reopen whatever we've closed, I think there's always a possibility for us to do that.

Okay. Thank you, Mark. Let's now go to the bulk transfer discussion based on Theo's input in the mailing list earlier this week. I think here we have two elements that are open. Let's start with Theo's one, which is 1E. Can a data escrow update mechanism be considered as part of a potential bulk transfer, an alternative option of bulk transfer, of existing registration data?

I understand, Theo, that some registrars are suggesting that the existing data escrow files could be used to some extent after some cleaning up to send the appropriate registration data to the registry.

Let me stop here and see if anybody would like to further discuss. I just want to mention that last week we had discussed that, if we are not using this as an alternative, and if we're closing these questions

EN

proposal [to] data escrow file, we would explain why we were not. So think positive or negative critical discussion of this point would be useful for us to record, and I will do so in our scorecard.

Let me stop here. Theo, please?

THEO GEURTS: Yes, thank you. Several registrars mentioned that, if they would use the RDE deposits, they would tackle the problem with missing data. I'm not sure how that works, but it was suggested. Several registrars suggested we just send one big file to Verisign, and they'd integrate it. That could an option. I'm not sure there. I think there may be a difficulty there when creating the context there that is maybe something Mark would comment on or maybe another IT member there. I'm not an expert there. Thanks.

FABIEN BETREMIEUX: Mark, please?

MARK MONITOR: Thank you, Fabien. I'll caveat this with saying I don't love this idea. It feels like it would be a not insignificant amount of work for us. But I'm not overly familiar with this RDE file. In fact, I don't think I'm familiar with the format or the contents of it.

I guess, if we want to evaluate it or consider it, if somebody could post a sample or maybe an example with some dummy data in it to the rest of the group, I could take a look at it and provide feedback. For example, one of the limitations we talked about is the [auth] information included

in that, for example. My sense is that would create more work in the end than it would actually save.

But, like I said, if somebody were to post an example of one of those documents, I would at least be willing to take a look at it and evaluate it and provide feedback to everybody else. Thank you.

FABIEN BETREMIEUX: Thanks, Mark. Would anybody on the registrar side be willing to share some data? Theo?

THEO GEURTS: Yeah. I'll look up the RDE deposit requirements from somewhere. So I'll look them up and pass them onto the list, and then we can discuss it further on the list, I guess. Thanks.

FABIEN BETREMIEUX: Great. Thanks, Theo. Any additional comments or questions on this subject? 1E: Data escrow update mechanism for bulk transfer.

Okay. If none, then we'll just leave 1E ongoing. We'll be looking forward to your input, Theo, in the discussion on the mailing list, and Mark's potential evaluation.

Okay. I want to go to 1A now. Before going to 1A, I just want to mention 1B, which we still have open, and that's the higher level discussion as to what alternative, if any, would be provided for the bulk transfer. I think Mark, in our last meeting, had explained what we would need to be considering in terms of activity in the bulk transfer.

EN

So I just want to mention it here, in 1B. If anybody would like to give some thoughts to this one, they would be more than welcome. Let me see if anybody would like to speak to 1B at all. If not, we'll go to 1A.

Okay. Not hearing anybody and not seeing any hands raised. So, 1A. We're back to the EPP standard way of providing the existing registration. I think we had 1A open because it was a discussion around whether there would be the capacity on the registries' side to handle high volumes of transaction with some of the registrars.

We had noted that, maybe, Mark, you'd have a sense by now, so let me just turn it over to you and see if you have any update in that area.

MARK MONITOR: Sure. Thank you, Fabien. I did look into this a bit with our technical teams. I guess there are a couple of considerations, the first being dedicated EPP connections to be made available. As Theo pointed out, we already have two channels for the .com and .net registries. We have a standard connection pool and what we call a batch connection pool.

Theo suggested that the batch connection pool would make sense for a pool for giving the bulk transfer. I took that back to our technical teams and they agreed. So what I would suggest, and what I think we would suggest as a policy, is that, for the backfill of contact information for existing registrations, we would direct registrars to use the batch EPP channel for all that uploading. I think this is to confirm what we talked about in previous sessions.

I saw Theo's hand go up. Did you want to jump in now?

EN

THEO GEURTS: No. You just finish, Mark. My question can wait.

MARK MONITOR: Okay. Fair enough. So the other thing was the volume or the level of transactions that we would be able to process for purposes of estimating traffic. Again, I looked at what we had done in the past for bulk updates to the system. What we've done in the past is we've requested that registrars limit their transactions to 30 tps per second, which, considering the amount of updates that need to happen, if you do the quick napkin math, is going to amount to months of time to backfill all the records. But considering we're talking about a 12 to 18-month window there, I think that's probably fairly reasonable.

An alternative to the 30 tps per second that was suggested by my technical teams was that, if registrars were to open up one channel and send as many transactions over that one channel as they could, they were confident that that would not cause issues with our system.

So if registrars to use multiple channels or multiple connections to us, we would ask that they limit it to 30 tps per second. But an alternative approach would be to limit the connections to one connection and then send as many transactions over that one connection as they were able to. Again, we'd ask that all the backfill, at least if you're using EPP, be limited to the batch connection pool.

I guess I'll stop there and see if there are any thoughts, comments, or feedback. Thank you.

EN

FABIEN BETREMIEUX: Thanks, Mark. Theo?

THEO GEURTS: Yeah. I just wonder, Mark, if this EPP batch pool would be an addition to and not be required? Is that correct? We haven't implemented that, so if we would we be required to use the back pool EPP, then it would mean development time, and [inaudible]. Thanks.

FABIEN BETREMIEUX: Go ahead, Mark.

MARK MONITOR: Yes, Theo. Sorry. I guess I'm asking you: so you don't use the batch pool at all? You currently only use the standard pool?

THEO GEURTS: That is correct.

MARK MONITOR: Okay. I guess we would ask that only the batch pool be used for the backfill of existing registrations if you're using the EPP approach for updating your contact information.

THEO GEURTS: But that would require us to implement that. Right?

MARK MONITOR:	Fair enough, but functionality-wise it's exactly the same. It's just different [inaudible].
THEO GEURTS:	Okay. Okay, I'll take this back to the team. Thanks so far.
FABIEN BETREMIEUX:	Thank you, Mark and Theo. Jody?
JODY KOLKER:	Mark, when you say 30 transactions per second, do you mean all transactions based on this, on replacing contacts? Or do you mean like a create contact, a domain update? I guess I'm trying to understand what the transactions per second is limited to. Is it just contact create?
MARK MONITOR:	Thanks, Jody. Yeah, I recognize that, in order to do the full backfill, it's not one transaction for every contact. It would be multiple commands. Basically, you would need to do multiple commands to backfill just one registration. I think that' what you're getting to there?
JODY KOLKER:	Right.
MARK MONITOR:	We wouldn't count it against your overall transactions. Certainly we don't want the backfill activity to impact any existing or standard

registration activity. But as far as the backfill activity, we would ask that you limit that to 30 tps.

JODY KOLKER: And that includes the contact create, the domain update, etc.? Right?

MARK MONITOR: Correct.

JODY KOLKER: Okay. All right. Thanks.

MARK MONITOR: So any initial thoughts on this? I did some napkin math on that against your registration, so I know that that amounts to a pretty significant undertaking for you. Any thoughts on that?

JODY KOLKER: Well, 30 transactions per second? Yeah, that's going to go pretty slow. But if we have 12 to 18 months to do this – well, I'm not sure if we're going to be able to get that done, actually. I guess we'll have to look at that. But I would echo Theo's concerns, though, too. If we're doing this for the batch pool, we'll have to do some development around that to do that to get it in there; to use the batch pool to do that instead of the regular pool.

And Theo was probably thinking the same thing, that in order to do this we would throttle this pretty slowly through the regular connections,

especially if it's going to be at 30 transactions per second. We could throttle it down pretty slow to go through the standard pool, especially if we're only going to use on connection to do it, or if we're going to be throttled at that rate. That's just my thoughts.

We were hoping to use the standard connections because there would be a minimum amount of code, and then we could throttle it slow enough that we would hope it wouldn't cause a problem with the registry system going through the standard pool like that.

I know we have the ability to be able to [inaudible] connections, but it would throttle it down so far that we would hope it wouldn't cause a problem.

MARK MONITOR: Okay. Thank you. I got to admit I was surprised. I didn't anticipate concerns with using the batch pool versus the standard pool. So let me take that back and see what my team thinks about that.

Also, I'm wondering, I haven't had a chance to get anybody to do tests on the number of transactions you would be able to do with one connection, but the thinking of the team is that, if you could limit the backfill activity to just one connection, that would protect our system but would also enable you to send more than 30 transactions per second. So I guess I'm wondering what your thoughts are on that option and if maybe you need more time to consider.

JODY KOLKER: I think I'm going to need a little more time to look at that, just using one connection. The way that I was planning on doing this through our system was that we would limit our system, basically, on the number of domains that it would be able to do it at once. So we wouldn't put 200 domains in there at one time and then open up all the connections. We would say, "All right. We're just going to work on one connection or one domain at a time and just run that through."

It might use a couple of different connections to do that, but it would just do one domain at a time. If that's what we need to throttle that, that's what we would do. If we can get away with two or three – we might use three connections – the way that we would feed it would be very slowly like that, and we wouldn't use just one connection. We might use four connections. But we would limit the number of domains that we would update at one time.

That would work well with our system. We would say, "You can do two domains in a minute," or three domains in a minute, but not 600 domains a minute.

MARK MONITOR: Gotcha. Okay. That's good feedback. Yeah, I'd like to take that back and ask the team to look over that and consider that some.

JODY KOLKER: Cool.

MARK MONITOR: Thank you.

EN

JODY KOLKER: Yeah. Thanks, Mark.

FABIEN BETREMIEUX: Theo, are you back in the queue?

THEO GEURTS: Yes, I'm back in the queue. It's a new hand. Mark, just a few more thoughts here. First is on the batch pool and the regular shared pool. It might be handy to look up how many registrars are actually using that batch pool versus the regular pool. That might be an idea. Also, combine that with the amount of registrations those registrars to have.

                         See, my feeling here is that, if I'm looking at Realtime Register here, the company I work for, we are not the size of a GoDaddy, so we are just assumed to go and do this through the regular shared connection there, the regular pool, because we were doing this. I feel like we got X domains here and we should be able to pump all the data over within a day over a regular connection in a worst case scenario. If we can make it any faster by using one connection or not, [that would] maybe just speed it up. So take [inaudible] if you can. Thanks.

MARK MONITOR: Yeah. Jody and Theo, thank you. That's real good feedback. I'd like a chance to review that internally and get back to you both on that.

THEO GEURTS: Yeah, and one more last thought here, Mark. Sorry. There might be also registrars out there that will require some significant planning if you would require the batch EPP pool. You can't just predict the flexibility of a registrar here. That's just what I'm saying. Thanks.

MARK MONITOR: Fair enough. I'm nodding. You can't tell, though.

THEO GUERTS: Okay.

FABIEN BETREMIEUX: Thanks, Mark. Yeah, that might be a future request: for Adobe Connect to have a nodding head as an icon.

Jody, I still see your hand. Is this a new hand? Okay. I assume it's an old hand. Thanks, Jody.

All right. So we'll leave 1A open. We'll be looking forward to Mark's review of that discussion and how that can move forward.

Let me now move to 7, which is the discussion of the timeline of the existing registration [track]. I wanted to get there because this is connected with the topics we're previously discussed. I believe, then, that 7A – it depends on the feedback you just provided, Mark, so I think that's helpful. We've discussed that and this will remain open until we can take that discussion further.

I want to mention 7B, which is the discussion around how the transfer by registrars would be organized or what type of incentives there would

EN

be to avoiding any problems during the time frame for backfilling the data or for the transition of the existing registration.

I'm curious here if we could maybe delve a little bit into some of the details here. I was reviewing the existing registration timeline, where we're talking about both the OT&E systems and the production systems going into an optional thick mode, and then a required thick mode.

I'm wondering if this was an opportunity to somehow, by requiring thick on the OT&E systems, push registrars to prepare for the transition to thick, and then be ready by the time the thick WHOIS is required on the production system. So I'm just thinking of this as a potential mechanism to somehow force the transition of everyone in time.

Let me stop here and see if anybody would like to discuss this topic further. Theo, you raised your hand?

THEO GEURTS: Thanks, Fabien. I've been giving this some thought, also, and I do not have a way forward here. When we were discussing the minimum points there on the registrar stakeholder list, I thought there was some good response there, and I thought there was some there. My feeling was that, for these guys from the Registrar Stakeholder Group – the registrars there – I don't think that is going to be the issue. We've got over 100 members there, so getting those guys organized? Just as an example, some of them are pretty eager to transfer over the data as fast as possible and be done with it. So you've got those kind of people.

So I don't think that the Registrar Stakeholder Group is the issue here. I think you need to expand your horizon there and look at the registrars

who are not a member of the Registrar Stakeholder Group because my feeling is, how do you go about with those people? I have no idea there. So this is my thinking here. Thanks.

FABIEN BETREMIEUX: Thank you, Theo. Mark, please?

MARK MONITOR: Thank you, Fabien. To answer your question, Fabien, about the OT&E, it's not a bad idea. To just give a little background, we typically see very low registrar activity in the OT&E environment. It's not particularly well-utilized on the whole. It's also important to note that OT&E doesn't have all the same registrations as production. So it's not like it's a mirror of production right there. They're two separate environments, so the many millions of registrations that exist in production don't necessarily exist in the OT&E environment.

So I'm not sure how effective that would be. It's certainly true that what we'd do in production would be done in OT&E first, so registrars would need to have their systems modified ahead of time for OT&E in order to use the OT&E environment. But it isn't well-utilized by registrars, so I don't know how effective that would be. Thank you.

FABIEN BETREMIEUX: Thanks, Mark. Thanks for sharing your feedback. Let me maybe follow up. You mentioned earlier the 12- to 18-month time frame the IRT had been discussing for the registration track. Is this something you're thinking that's also a reasonable time frame for the backfill of data,

independently of the discussion on the technical parameters that may impact the [inaudible], etc.? Is this something you're thinking as a ballpark figure that we can start working with as a starting assumption?

MARK MONITOR: I think I'd really have to defer to registrars for that. I think it's much more a question for them than it is for the registries. So I'm going to try to completely dodge that one, Fabien. Thank you.

FABIEN BETREMIEUX: Thank you, Mark. Anybody interested to share a sense of what they're thinking on 12 to 18 months as a window for backfilling the data? I think we're trying to get to place where we start with numbers on the timeline and start drawing a timeline over some kind of calendar and try to get a sense of where we're going.

Let me stop here. Jody? Are you [there]?

JODY KOLKER: I was. Sorry. I'm looking at how long this would take us at 30 transactions per second and how many transactions we think we need. It's going to take, I want to say, at least six or seven months, and that's if we don't have any issues. And that's going pretty – I don't want to say pretty consistently that long. So I think 12 to 18 months, because I'm expecting to have issues, as for some reason the service might quit. We might have to restart it and make sure we're keeping an eye on that not happening. I'd be interested to hear anybody else's thoughts, though.

I would say at minimum at least 12 months, and I'd feel more comfortable with 18.

FABIEN BETREMIEUX: Thank you, Jody. Thanks for sharing that perspective. Theo?

THEO GEURTS: Yeah. Just to circle back around to Jody there, I think the request for 18 months is not unreasonable. If the timeline is so much correct there, like 8 to 9 months to do it, hopefully it is best case scenario there and nothing goes wrong. But assuming stuff goes wrong there, then already Jody and his team would be in a bad spot there. So I don't think the 18 month request is unreasonable. Thank you.

FABIEN BETREMIEUX: Thank you, Theo. Steve?

STEVE METALITZ: Yes. Thank you. Just to make sure I understand how this fits in with the earlier conversations, is this the time dictated by the limitation of 30 transactions per second and the reluctance of registrars to use this batch method that was being talked about? Or am I mixing apples and orange here?

FABIEN BETREMIEUX: Thanks for your question, Steve. Jody? Theo? Mark? Do you want to share your thoughts?

JODY KOLKER: Sure. I'm looking at, for every domain – and we haven't completely thought about how to do this yet, but it's going to be, at minimum, five transactions a call for each domain name. I got a feeling that that could be as high as ten transactions. So if I do ten transactions on 62 million names, that's going to be 620 million transactions times 30 per second. I'm coming up with between six and twelve months, and that's if I don't find anything else that we need to do there.

Definitely, we have not done much back-of-the-envelope here in thinking about how we're going to do this. There are probably efficiencies that could be done. The more you dig into it, the deeper you get in, or the more things you find, as in "This doesn't work right," or, "We should be doing it this way." So I'm trying to lead a [path] in there so that we don't miss whatever date we put up. That's my concerns with that.

FABIEN BETREMIEUX: Steve, would you like to respond?

STEVE METALITZ: No. I think I understand that. So I understood the earlier conversations. There was also a method that could be used that wouldn't be subject to the 30 transactions per second. Are you ruling out the use of that method? Or again am I confusing two different things?

JODY KOLKER: Was that meant for me or Theo?

THEO GEURTS: No, I can answer it, Jody. I don't mind.

JODY KOLKER: Okay.

THEO GEURTS: Steve, we are not ruling that method, that alternative path there, where we just upload it in one big file. But as Mark mentioned earlier on on the call, we need to know what is actually required. I'm going to send the information to the list later this week, so Mark [inaudible] can have a look at it. If it's a usable, then we can discuss it further. Maybe that is a path there to speed things up. So that is maybe a path forward there. Thanks.

STEVE METALITZ: Thank you very much.

FABIEN BETREMIEUX: Thank you. Jody? Please, go ahead.

JODY KOLKER: Sorry, Fabien. Yeah, from our perspective, we weren't planning on using the file method at all. We were just going to go through this as part of our regular processing.

To me, I feel that there's going to be a lot of back of forth using the file method between the registry and registrar. We're trying to avoid that as

much as possible because, if you find problems with the file, then you have to send it back to Verisign, and there could be a lot of back and forth there.

As far as we're concerned as GoDaddy, we're going to trickle this in as fast as Verisign will let us do that or as fast as the registries will let us do that, because we're hoping for is a fire-and-forget, basically. We'll look at the failures and we'll fix those, and then we'll fix them going forward, I guess. I'm hoping this will be the least amount of time on developer time, both at the registry and the registrar. Just my thoughts.

FABIEN BETREMIEUX: Thank you, Jody. Theo, I see you in the queue. Is this a new hand?

THEO GEURTS: Yeah, this is a new hand. Yeah, I'm just going to repeat this, what Jody just said. Realtime Registry's developers are looking at using the RDE deposits also because they also feel there is going to be some back and forth there. There could be some data corruption there. There are some technical limitations to it. So it's not a Holy Grail there, so to speak. But there were some registrars who pointed it out, and I think as the IRT should be definitely looking at and see if it's a possible way forward.

But, yeah, we had the same reservations as Jody just pointed out there. Thank you.

FABIEN BETREMIEUX: Thank you, Theo, for sharing that feedback, and thank you, Steve, for bringing this to the fore. It seems that, in order to move forward with a

EN

sense of timeline without knowing what the available options are and what this means to the various types of actors, I think we've identified a few options on this call which will certainly help us move forward to a place where we can set a timeline.

Steve, please?

STEVE METALITZ: Thanks. I'll just reflect back what I'm hearing, which is that there might be a method that could be faster but the registrars are reluctant to do it because there could be complications with it. That leaves the option in which Verisign sets the pace, and they've decided it's 30 transactions per second. I'm not quite sure that is and whether that could be subject to revision in a way that would enable a faster completion of this job.

That's how I understand the lay of the land right now. Obviously, I may well be wrong about that, and I would welcome any corrections.

FABIEN BETREMIEUX: Thank you, Steve. I see Roger in the queue next.

ROGER CARNEY: Hi. Just to respond to Steve, I don't think that we're saying one is faster. We're saying it could be faster for some people or larger registrars. It's not going to be a simple one file in which the registry will solve the problem. You're not going to send 160 million records in one file and get it back. So I don't think that it's faster. It may be faster for some people, for some registrars, but not for the larger registrars. Thank you.

FABIEN BETREMIEUX: Thank you, Roger. Theo?

THEO GEURTS: Thank you, Fabien. Steve, just to highlight this, we are discussing it, and like Roger just actually said, this might be faster. It could also be slower. I think the bottom line here is or the key item here is that we've never done this before, so we are actually discussing something we have not much experience with, so to speak. Thanks.

FABIEN BETREMIEUX: Okay. Thank you, Theo. Roger, I see you in the queue still. Is this a new hand or old hand?

ROGER CARNEY: Sorry. Old hand.

FABIEN BETREMIEUX: Okay. Thank you. Steve, yes. Please, go ahead.

STEVE METALITZ: Just in response to the last comment, it has been done before. We've done it in .org. That wasn't obviously on the same scale, but it was large. So I wonder whether are learnings from that that could be adapted.

                      But I hear what you're saying. I understand that this other option is an unknown at this point. So it could be faster or it might not, and that

brings it back to the question of why the first option is set at 30 transactions per second. Thank you.

FABIEN BETREMIEUX: Theo?

THEO GUERTS: Just to go back real quick there to Steve's comments there, the alternative of the RDE deposit wasn't available back in the .org migration because we didn't do RDE deposits. We started the RDE deposits after the registry [inaudible] disaster, which was in 2007, I think – well, at least the [RDE] was launched there. So it's out there now as a possibility, but we can't put any timelines on if it will be better, faster, or longer. Thank you.

STEVE METALITZ: Thank you.

FABIEN BETREMIEUX: Jody?

JODY KOLKER: I can remember the pain of moving .org from thin to thick. I believe we only had about 600,000 domains at that time, and .org was maybe two or three million at the time. Somebody can correct my numbers there for .org, but I think it was definitely under ten million. Now we're talking about a domain that's at least ten times as large as that. For us, it's 100

EN

times as large. This was quite painful for the 600,000 domains that we had to do.

The painful part of it wasn't just the transfer of it. It was after we were supposedly done dealing with the registrars that hadn't moved their data over to the thick registry. We were constantly getting pinged by Afilias or PIR that the data that we had in our database was incorrect, although they were transfers. They were transfers of domains from other registrars, and I believe we were getting that years after the transfer was done.

So I guess I can communicate the pain of doing it. That's why [I'd like] to do this very carefully as we're thinking about it. I guess what I want to say is that I think that Theo has communicated the right position that the members of the Registrar Stakeholder Group that are at meetings are not going to be the problem here. It's going to be the registrars that maybe have a few thousand domain names that haven't translated their domains yet and wow we handle that. Is there supposed to be repercussions for that? If they haven't moved over in 12 or 18 months, do you stop allowing them to register domain names? How does that work?

FABIEN BETREMIEUX: Steve?

STEVE METALITZ: Well, I could just give you the legal answer to that, which is that they have to comply with ICANN consensus policies. This is a consensus

EN

policy. So, yes, if there's no other way to do it, then perhaps you have those consequences.

But I hear what you're saying, that it's the registrars that are already active and engaged that are most likely to do this most quickly and that there will always be outliers. I'm sure that's true, but I think there are at least some legal tools to deal with that. Thanks.

FABIEN BETREMIEUX: Thank you, Steve, and thank you, Jody. I'm checking the timeline. We have two minutes less, so unless somebody would like to close this discussion – I've taken a number of notes here on as many elements that we need to push through the discussion on so that we get to place where we believe we have a solid plan. So this is very positive from my perspective, I think.

Joyce? Yes, please. Joyce, are you muted? Joyce, we cannot hear you.

All right, so Joyce, please, if you're able to speak, let me know. In the meantime, I just want to remind everyone that our next IRT meeting will be next Tuesday. Before Helsinki and during Helsinki, we have two slots. The currently the time on the agenda on those ones, so we're not necessarily [inaudible] so. We'll have to confirm, but this is currently what's on the schedule.

Thank you again for your time today. We'll be looking forward to the extension on the mailing list, and we'll try to share the updated scorecard as soon as possible so it can help with the discussion.

Thank you, Roger, and all, for your congratulations. I was very happy to be with you today and we're very happy to [inaudible]. Thank you all.

[END OF TRANSCRIPTION]